arXiv:1701.04905v1 [q-bio.NC] 18 Jan 2017

Feedforward Architectures Driven by Inhibitory Interactions

Yazan N. Billeh[1,*] and Michael T. Schaub[2,†]

[1] *Computation and Neural Systems Program, California Institute of Technology*[‡]
[2] *ICTEAM, Université catholique de Louvain*[§]

Directed information transmission is a paramount requirement for many social, physical, and biological systems. For neural systems, scientists have studied this problem under the paradigm of feedforward networks for decades. In most models of feedforward networks, activity is exclusively driven by excitatory neurons and the wiring patterns between them while inhibitory neurons play only a stabilizing role for the network dynamics. Motivated by recent experimental discoveries of hippocampal circuitry and the diversity of inhibitory neurons throughout the brain, here we illustrate that one can construct such networks even if the connectivity between the excitatory units in the system remains random. This is achieved by endowing inhibitory nodes with a more active role in the network. Our findings demonstrate that feedforward activity can be caused by a much broader network-architectural basis than often assumed.

I. INTRODUCTION

The ability to reliably propagate signals in a targeted manner is essential for the operation of many natural systems, and a necessary building block to establish many further computational mechanisms. Prototypical models for such targeted information transmission within a neural substrate are feedforward networks, which have been considered in the literature for decades. In these models, the basic paradigm is to group nodes (neurons) into separate layers, each of which receives excitatory input from the preceding layer, and projects excitatory connections to the subsequent layer (see reviews, [1, 2]). The thus established forward-directed excitatory pathways guide the activity sequentially through the layers. A large number of variations of this scheme have been considered, such as embedding feedforward architectures in randomly connected networks to examine their effect on the overall network dynamics and signal propagation [3, 4]. Further, feedforward networks have been shown to propagate firing rates [5, 6], synchrony/pulse packets [7–12], combinations of firing rates and synchronous spiking [2], and even the ability to gate activity transmission [13]. However, within the paradigm of feedforward networks, inhibitory units (neurons) play merely a balancing role: they ensure that the network remains stable, either separately for every layer or globally.

Interestingly, recent work in neuroanatomy has revealed an enormous diversity of inhibitory neurons [14–22]. Moreover, specific plasticity rules for different subtypes of inhibitory neurons [23] add further to their diverse and heterogeneous connectivity profiles. In this light it would be strongly surprising if inhibitory neurons serve the cortex only in a homogeneous, passive role when it comes to information propagation, as is assumed in most feedforward networks. While some studies exist that specifically account for intra- and interlayer inhibition [3, 6, 24–26], information is still mediated by a cascade of excitatory neurons. However, as experimental evidence suggests, the nervous system likely uses a combination of methods to transmit information [27, 28]. An important natural question is thus if embedded excitatory pathways are necessary for feedforward propagation or whether one can construct networks in which there are no preferred excitatory-to-excitatory pathways, but *inhibitory neurons* play a pivotal role for the propagation of activity between layers. In the following we demonstrate, via numerical experiments and brief analytical considerations, that such feedforward processing is indeed possible with two exemplary circuits.

II. RESULTS

A. Cross-coupled feedforward networks

In a recent experimental study of the hippocampus [29] it was demonstrated that excitatory neurons in the CA3 region are unable to drive excitatory neurons in area CA2 due to strong feedforward inhibition. However, when this strong inhibition of CA2 is alleviated, CA3 can indeed excite CA2 excitatory cells to elicit action potentials [29] (Fig. 1a). An interesting feature of this finding is that the directionality in the interaction between CA2 and CA3 appears to be dictated by the connections between excitatory and inhibitory neurons, rather than a consequence of unidirectional excitatory connections targeting CA2. Indeed, excitatory connections between CA2 and CA3 are reciprocal [30, 31] – yet there is still a *directed* propagation towards CA2 [29]. Stated differently, the targeted activation of CA3 is controlled by an excitatory-inhibitory-excitatory pathway.

We sought to leverage this targeted activation mechanism by cascading this connection motif using leaky-integrate-and-fire (LIF) networks (see Materials and Methods). The result is a circuit with *uniform* connectivity among excitatory neurons with *no preferred direc-*

* yazanb@alleninstitute.org
† mschaub@mit.edu
‡ *Current Address:* Allen Institute for Brain Science, Seattle
§ *Current Address:* Institute for Data, Systems, and Society, Massachusetts Institute of Technology

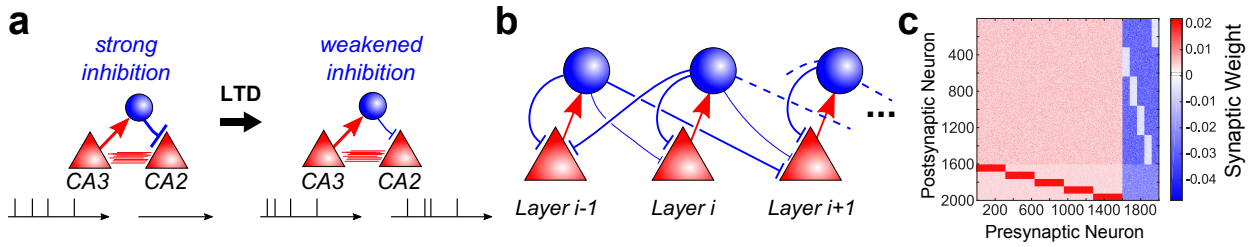

Figure 1. **Cross coupled feedforward networks** . **(a)** Schematic of the finding of Ref. [29]. Inhibitory long term depresssion (LTD) lowers the feedforward inhibition of CA2, allowing information transfer from CA3 to CA2. **(b)** Schematic of proposed network architecture. Note that only connections that are not identically distributed in the rest of the network are displayed for visual clarity. The feedback between excitatory and inhibitory neurons drive the feedforward activity, while connections between alike neurons remain uniform (see text). **(c)** Synaptic weight matrix of an example network with five groups. Note that the network contains connections between all neuron types. Panel **b** only emphasizes the dominant pathways altered in the ccFFN architecture.

tion, which is nevertheless able to propagate feedforward activity due to the specific cross coupling between the excitatory and inhibitory neurons. We term this circuitry a cross-coupled feedforward network (ccFFN). A schematic can be found in Fig. 1b that emphasizes the main properties of the network only and does not include all connections. For that, Fig. 1c is a weight matrix of a network instantiation with 5 groups (with the last group connecting to the first) which illustrates the full connectivity of ccFFNs.

The behavior of ccFFNs can be explained by the following rationale: (i) the excitatory neurons in each layer are more strongly coupled to the group of inhibitory neurons in their own layer relative to other inhibitory neurons; (ii) an activity increase of such an excitatory group thus triggers elevated activity in the corresponding inhibitory neurons; (iii) this inhibitory group of neurons targets the subsequent layer of excitatory neurons more weakly relative to other excitatory neurons; (iv) the reduced inhibition (relative) of the subsequent excitatory group leads to increased excitatory activity in the subsequent layer, while the activity in the initial layer returns to baseline; (v) by cascading this cross-coupling motif, elevated activity of excitatory neuron groups propagates through the circuit.

A simulation of a network with such a ccFFN topology is shown in Fig. 2a. Note that, to eliminate transient effects from particular driving inputs we connected the last layer with the first layer, thus establishing a circular pathway with a self-sustained forward propagation of activity. Importantly, in addition to the propagation of the excitatory activity, we observe that the *inhibitory* neurons' activity progresses from one group to the next. This emphasizes the pivotal role played by inhibitory units for the observed dynamics, which is clearly beyond simply balancing the network.

For our simulations, to describe the statistical strength of the aforementioned connectivity motifs, we have defined a 'forward activity' parameter Q, which is simply the ratio of the connection probabilities of the 'targeted' vs. 'non-targeted' neuron groups (or the ratio of corresponding synaptic weights, respectively), that modulates the amount of feedforward structure (see Materials and Methods). For simplicity, we kept all these ratios equal. Note, however, that feedforward activity can be observed by changing only the weights or the connectivity probabilities separately (for a related observation, see [32]). Using this simple setting, by varying Q as our only parameter, we can alter the overall feedforward structure. Note, $Q = 1$ corresponds to the case where the network is perfectly uniform. By increasing Q to values larger than 1 the level of feedforward structure increases.

The network simulated in Fig. 2a consists of 5 layers of neurons with a feedforward ratio of $Q = 2.6$. To illustrate that increasing Q indeed results in increased feedforward activity, we calculated the cross-covariance (Fig. 2b) and the Pearson correlation coefficient (Fig. 2c), between the firing patterns of the neurons, averaged over the different layers.

Fig. 2b shows the average cross-covariance functions within the same layer for the two conditions of $Q = 1$ and $Q = 2.6$, averaged over 10 realization of the network. To get a smooth estimate, we convolved the spike-train of every neuron with a Gaussian signal of standard deviation $5ms$. For every neuron pair, the convolved signal ($f_i(t)$) was then used to calculate the pairwise cross-covariance $\phi_{ij} = \mathrm{cov}[f_i(t+\tau), f_j(t)]$:

$$\phi_{ij}(\tau) \approx \int [f_i(t+\tau) - \mu(f_i)][f_j(t) - \mu(f_j)]dt, \qquad (1)$$

which we averaged over all neurons inside the same layer: Here $\mu(\cdot)$ denotes the mean of the signal. While there is no apparent temporal structure in the networks with $Q = 1$, there is a clearly visible increased synchrony in the networks with high feedforward ratio $Q = 2.6$, as indicated by the large peak at zero lag. Moreover peaks appearing at a lag of $\pm\delta$ corresponding to the repetition period of the firing, resulting from the the circular topology we imposed.

To investigate the tendency for each layer to fire in unison further, we computed the Pearson correlation coefficient of the convolved spike-trains of all neuron pairs

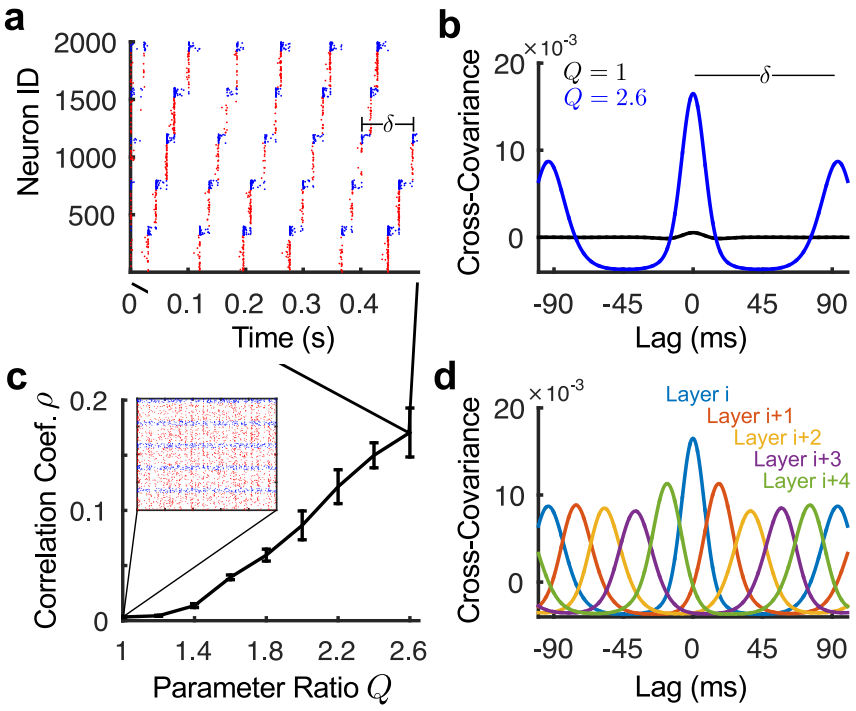

Figure 2. **Firing pattern characterization of cross coupled feedforward networks** . **(a)** Example raster plot with 5 groups that show the propagation of activity between layers. Observe that inhibitory neurons also show feedforward propagating activity and that the final group connects back to the first (circular arrangement) and hence the activity propagates indefinitely **(b)** Group-averaged cross-covariance for parameter ratio values $Q = 1$ and $Q = 2.6$. For $Q = 1$ (no feedforward structure) the firing is clearly not synchonous and does not display any pattern. In contrast, with imposed cross-coupled feedforward structure ($Q = 2.6$), there is a peak indicating strong synchronous firing inside each layer. The second peak at time $\pm\delta$ indicates the periodically repeating firing pattern (see also **(c)**). **(c)** The average Pearson correlation coefficient within layers as a function of Q. The larger the feedforward ratio, Q, the greater the correlation of firing within layers. Error bars are standard deviations. **(d)** Average cross-covariance of the neural firing patterns in layer i with neurons in all other layers j for $Q = 2.6$. Observe that the time lag of the peaks are arranged consecutively, illustrating the orderly feedforward progression between groups.

for varying levels of Q. We plot the average correlation coefficient within each group and layer in Fig. 2c. For each value of Q, 10 network realizations were simulated. As can be seen, the larger the value of Q, the more synchrony there is within groups. This synchronous firing of groups is not decoupled but propagates along layers as can be seen in Fig. 2d. There we plot the average cross-covariance of a layer relative to all other layers (c.f Fig. 2b). As the regular shifts in the cross-covariance indicate, there is indeed a clear consecutive progression of activity from one layer to the next. We can thus conclude that the ratio Q directly influences the feedforward activity propagation in ccFFNs, demonstrating that directed information transmission is possible without an imposed excitatory-to-excitatory pathway in the network.

B. Disinhibitory feedforward networks

The aforementioned cross-coupling of excitatory and inhibitory neurons is not the only arrangement possible to create targeted feedforward activity driven by inhibitory units. To illustrate that our finding is general and should not be reduced to a single type of circuitry, we present here a second network architecture which takes inspiration from the specificity of inhibitory neurons' connectivity patterns [17, 18]. In disinhibition motifs, certain subtypes of interneurons inhibit other interneurons which normally suppress connected excitatory neurons. Such a disinhibition cascade can thus lead to an increase in firing rates in excitatory neurons which are normally suppressed. Interestingly, long range disinhibition has been reported to propagate network activity [33, 34], although we stress the architecture presented here does not impose spatial constraints.

Inspired by these experimental findings, we constructed a network model in which disinhibitory motifs enable the feedforward propagation of spiking activity. We hereafter call this architecture a disinhibitory feedforward network (dFFN). A diagram of the wiring scheme for a dFFN is shown in Fig. 3a that highlights the key features of the network and does not include all connections. For that, Fig. 3b shows a weight matrix instantiation with 5 groups (with the last group connecting to the first) which illustrates the full connectivity of dFFNs.

The functionality of this circuit can be explained

schematically as follows: (i) each layer comprises a functional group of excitatory and inhibitory neurons more strongly connected to each other than to the rest of the network; (ii) inhibitory neurons in one layer target preferentially the inhibitory neurons in the subsequent layer (disinhibition); (iii) thus, when the activity in the preceding layer increases, the inhibitory neurons' activity in the next layer will decrease. (iv) This in turn allows the excitatory neurons in the next layer to increase their firing; (v) after a short delay, the inhibitory neurons increase in activity again, as they receive input from the excitatory neurons in their layer, which have elevated activity as a result of their disinhibition. This eventually reduces the total firing back to baseline in the group – however, not without the activity moving to the next group via disinhibition again; (vi) by cascading this motif, every upstream layer is activated and the information propagates. We implemented dFFNs with varying number of layers confirming that the above circuitry results in directed information propagation (Fig. 3a). Once again we observe feedforward signal propagation for both the excitatory and inhibitory neurons (Fig. 3c). As before, to avoid boundary effects, we used a circular network layout in which the last group connects back to the first and hence the activity keeps propagating.

Similar to the ccFFNs, the level of feedforward structure was control by a parameter Q that determines the ratio of connection probabilities and ratio of weights in the network to realize a dFFN (see Materials and Methods). The displayed network in Fig. 3c corresponds to a 5-layer network with $Q = 2.6$. When comparing the average cross-covariance function of such a network with uniformly connected networks $Q = 1$, we can again see an increased synchrony and a propagation of activity resulting from the dFFN architecture (Fig. 3d). This is further demonstrated by the increasing Pearson correlation coefficient of neurons within layers as a function of Q (Fig. 3e). Finally, we plot again the average cross-covariance between different layers in Fig. 3f. Although the same behavior is qualitatively seen of the signal propagating along the cascaded layers of the dFFN topology, this appears to be far less pronounced than in the ccFFN architecture. The reason for this effect can be explained by inspecting the raster plot further (Fig. 3c). Note that, in the ordered rastergram, the lowest group can fire again even though the cascade that emanated from it previously has not reached the top group yet. For instance, the firing of a new cascade at layer i and a previous cascade at layer $i + 3$ can temporally overlap and this co-alignment results in the multiple peaks in the cross-covariance. Hence, the network propagates multiple signals concurrently, or stated differently, the network's architecture is able to process multiple signals simultaneously effecting the cross-covariance measure.

C. A stylized linear rate model for disinhibitory feedforward networks

In order to gain some further insight into the operation of the disinhibitory feedforward network, it is insightful to consider a stylized linear rate model [32]. Let us denote the group of excitatory (inhibitory) neurons in layer 1 by E_1 (I_1), and the corresponding firing rates by r_{E_1} (r_{I_1}); see Figure 4a for a schematic. Note that without loss of generality we may assume that the firing rates are measured relative to some baseline activity, and can thus be positive or negative. We can then write down the following coupled rate equations for layer 1 (equivalent equations for layer 2 not shown):

$$\frac{dr_{E_1}}{dt} = -\beta_{E_1} r_{E_1} - r_{I_1}, \tag{2}$$

$$\frac{dr_{I_1}}{dt} = -\beta_{I_1} r_{I_1} + r_{E_1} - r_{I_2}, \tag{3}$$

where β denotes the effective self-coupling of the individual groups, which is a compound of a constant decay (leak) term and the coupling of the neurons within the same group. For simplicity we assume here that there is no self-excitation term for the excitatory neurons, and hence the network has an excess of inhibition (not balanced). The magnitude of the parameter β will control how much dissipation there is for each neuron. If we set $\beta = 0.5$, which is the minimal requirement for a stable system, the dominant modes are associated with purely imaginary eigenvalues and correspond to (phase-shifted) wave-functions.

If the network is now excited by an instantaneous input, such as an input to E_1 at $t = 10$ as shown in Figure 4b, indeed a propagation of activity can occur along the network as outlined above. Note that the general features of this finding are not an artifact of the chosen parameters or model size, but can be generalized to arbitrarily sized networks. The forward propagating activity is indeed a consequence of the described architecture: the particular inhibitory connection profiles drive the dynamics.

III. DISCUSSION

We have demonstrated, via LIF network simulations, that one can construct networks in which feedforward activity is propagated even if connections between excitatory neurons are kept completely random. This is achieved by endowing the inhibitory neurons with an active role in the feedforward propagation, which has not been considered so far in the literature. Two circuit layouts were proposed demonstrating that both inhibitory an excitatory neurons can display feedforward dynamics simultaneously. We remark that the 'layers' discussed within both of these network types should not be taken too literally. While within our wiring schemes certain subpopulations are targeted (statistically), there is still a (biased) all-to-all connectivity probability

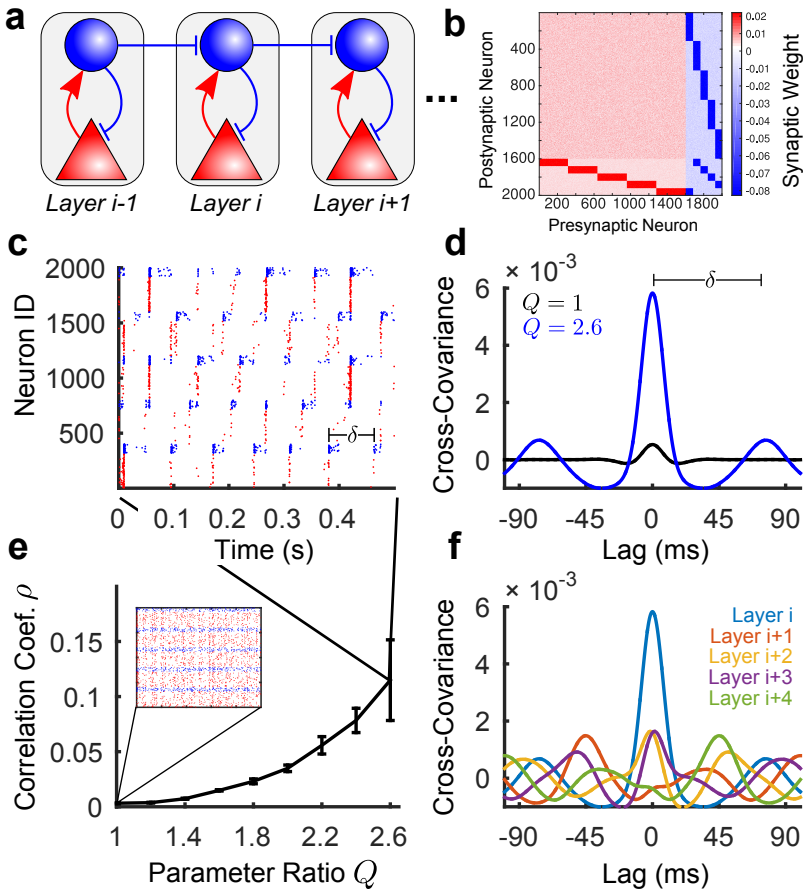

Figure 3. **Disinhibitory feedforward network. (a)** Schematic of dFFN architecture, in which feedforward activity propagates by the cascading of several disinhibitory structural motifs. For clarity, not all connections are shown in the schematic; importantly, excitatory to excitatory connections are randomly connected (see text). **(b)** Synaptic weight matrix of an example network with five groups. Note that there are connections between any types of neurons, as described in the Method section (**a** only emphasizes the pathways altered in the dFFN architecture). **(c)** Example raster plot of a network with 5 layers showing forward propagation ($Q = 2.6$). The inhibitory neurons again display forward propagating activity. Note that the final group connects back to the first (circular arrangement) and hence the activity propagates indefinitely. **(d)** Cross-covariance functions for the cases of $Q = 1$ and $Q = 2.6$ (see Fig. 2). **(e)** The average Pearson correlation coefficient within layers as a function of Q. The larger the feedforward ratio, Q, the greater the correlation of firing within layers. Error bars are standard deviation. **(f)** Average cross-covariance function between different layers (see also Fig. 2). Note that due concurrently propagating multiple cascades, the crosscovaraince between different groups display multiple peaks (see text).

throughout. Circuits like discussed here might thus be found within one cortical column, for instance. While long range projections in mammalian brains tend to be excitatory, topologies as discussed here might also be implemented by long range excitatory connections targeting local inhibitory neurons (see Fig. 1a). Moreover, long-range inhibitory projections have been reported [17, 34–42]. However, we emphasize that our models are not dependent on projection lengths and thus are not limited to a particular scenario by construction.

The first circuit was inspired by the hippocampal architecture recently uncovered between CA3 and CA2 where the propagation of a signal between CA3 to CA2 pyramidal cells is governed by the amount of inhibition CA2 interneurons impose on their CA2 pyramidal cells [29]. Interestingly, there exist some further experimental indications that a similar mechanism to direct activity might be implemented in canonical cortical microcircuits [43]. The traditional view of these ubiquitous circuits is that inputs from layer 4 (L4) drive layers 2/3 (L2/3) that then excites layer 5 (L5). However, as Pluta and coworkers [43] have shown, the picture is likely to be more intricate: in particular it appears that L4 first suppresses L5 while driving L2/3, and only afterwards L5 shows elevated activity [43] — a finding that shows parallels to our proposed ccFFN mechanism.

The second circuit was inspired by the disinhibitory role of interneuron subtypes [17, 18, 44] in addition to the subtype specific plasticity rules [23]. Such discoveries and the advancement of connectomics in uncovering the diversity of neuronal cell types and their corresponding connectivity rules influenced our proposed wiring scheme

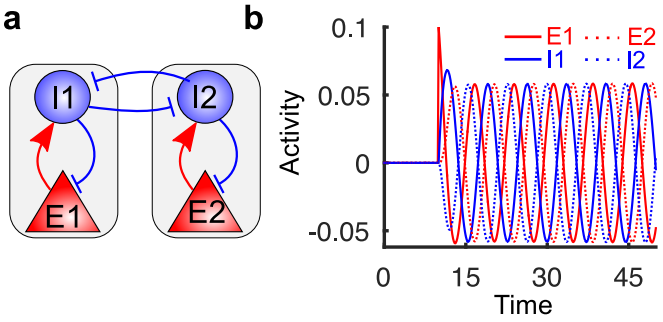

Figure 4. **Inhibitory connection profiles drive dynamics. (a)** Network architecture of stylized linear model. See text for dynamical equations. **(b)** Example simulation for $\beta = 0.5$ and a pulse of activity given to E_1 at $t = 10$ causing a permanent oscillation in all nodes of the network.

of dFFNs.

Overall, the recently discovered diversity of interneurons suggests that they play a much more vital role for the dynamics than simply acting as a balancing device for the network. We believe that new experimental discoveries call for a reassessment of the role of inhibitory neurons in models or neuronal circuits, and would encourage scholars to assign them more active and functional roles. Here we have demonstrated how such a functional role could be shaped in feedforward networks, but possibilities for such roles could clearly go far beyond.

IV. MATERIALS AND METHODS

Simulations were performed in MATLAB (2012b or later) and code can be found at `github.com/CellAssembly/inhibitory-feedforward`.

A. Model of spiking neurons

To illustrate our ideas, we have used leaky-integrate-and-fire (LIF) networks, stylized models of neural networks, which act like pulse-coupled oscillators. Using a time step of 0.1ms we numerically integrated the non-dimensionalized membrane potential of each neuron, which evolved according to:

$$\frac{dV_i(t)}{dt} = \frac{1}{\tau_m}(\mu_i - V_i(t)) + \sum_j W_{ij} g_j^{E/I}(t), \qquad (4)$$

with a firing threshold of 1 and a reset potential of 0. All networks comprised $N = 2000$ units, with an excitatory to inhibitory neuron ratio of $4:1$ (1600 excitatory, 400 inhibitory). The input terms μ_i were chosen uniformly in the interval $[1.1, 1.2]$ for excitatory neurons, and in the interval $[1, 1.05]$ for inhibitory neurons. Membrane time constants for excitatory and inhibitory neurons were set to $\tau_m = 15$ ms and $\tau_m = 10$ ms, respectively, and the refractory period was fixed at 5 ms for both excitatory and inhibitory neurons. Note that although the constant input term is supra-threshold, balanced inputs guaranteed an average sub-threshold membrane potential [32, 45].

In the model, the network coupling is captured by the sum in (4), which describes the input to neuron i from all other neurons in the network. Here W_{ij} denotes the weight of the connection from neuron j to neuron i ($W_{ij} = 0$ if there is no connection). After a presynaptic spike of neuron j, the synaptic inputs $g_j^{E/I}(t)$ are increased step-wise ($g_j^{E/I} \to g_j^{E/I} + 1$) instantaneously, and then decay exponentially according to:

$$\tau_{E/I} \frac{dg_j^{E/I}}{dt} = -g_j^{E/I}(t), \qquad (5)$$

with time constants $\tau_E = 3$ ms for an excitatory interaction, and $\tau_I = 2$ ms if the presynaptic neuron is inhibitory. For all networks described in the following, the total connection-strength per neuron was kept equivalent to an unstructured, balanced network displaying asynchronous activity, with $p_{EI} = p_{IE} = p_{II} = 0.5$, $p_{EE} = 0.2$, $w_{EI} = w_{II} = -0.042$, $w_{IE} = 0.0115$, and $w_{EE} = 0.022$. Here, p and w stand for the connection probability and connection weight, respectively. The first subscript denotes the destination and the second superscript denotes the origin of the synaptic connection, and E,I stand for an excitatory or inhibitory neuron, respectively.

B. Cross-coupled feedforward networks (ccFFN)

To construct ccFFNs we kept excitatory-to-excitatory and inhibitory-to-inhibitory connections uniform as outlined above. We divide the network into layers consisting of both inhibitory and excitatory units and connected as outlined in Fig. 1b,c: Excitatory neurons in layer i are statistically biased to target the inhibitory neurons in their own layer with a weight ratio $W_{IE} = w_{IE}^{i}/w_{IE}^{not[i]}$, compared to the inhibitory neurons in the rest of the network. Similarly, the inhibitory neurons within a layer i target the excitatory neurons in the *next* layer $i+1$, more weakly according to the ratio $W_{EI} = (w_{EI}^{i+1}/w_{EI}^{not[i+1]})^{-1}$. In addition to modifying the weights, we control the analogous ratio of connections probabilities R_{IE} and R_{EI}. Note that W_{EI}, R_{EI} are defined with an inverse ratio, i.e., a higher ratio means a *weaker* targeting corresponding to a stronger feedforward structure. To modulate the embedded feedforward level in the networks, we can thus vary the ratios $W_{IE}, W_{EI}, R_{IE}, R_{EI}$, while keeping the average weights and number of connections constant.

C. Networks driven by disinhibitory structure (dFFN)

In dFFN networks, excitatory-to-excitatory connections remain again uniform. Every layer in this network is composed of groups of exctiatory and inhibitory units as shown in Fig. 3a,b. Similar to above, we can control the imposed feedforward level with the ratio parameters $W_{IE} = w^i_{IE}/w^{not[i]}_{IE}$, $W_{EI} = w^i_{EI}/w^{not[i]}_{EI}$, and $W_{II} = w^{i+1}_{II}/w^{not[i+1]}_{II}$ or the analogous ratios of connections probabilities R_{IE}, R_{EI}, R_{II}. Again, for simplicity we set all six parameters equal to Q and vary them concurrently.

ACKNOWLEDGMENTS

We are thankful for discussions with Christof Koch, Costas Anastassiou, and Mauricio Barahona and comments from Jean-Charles Delvenne and Renaud Lambiotte. YNB wishes to thank the Allen Institute founders, P. G. Allen and J. Allen, for their vision, encouragement and support. Most of this work has been performed while MTS was at the Université catholique de Louvain. MTS acknowledges support from the ARC and the Belgium network DYSCO (Dynamical Systems, Control and Optimisation) and an F.S.R. fellowship of the Université catholique de Louvain.

[1] T. Vogels, K. Rajan, and L. Abbott, Annu Rev Neurosci Annual Review of Neuroscience, **28**, 357 (2005).
[2] A. Kumar, S. Rotter, and A. Aertsen, Nat Rev Neurosci **11**, 615 (2010).
[3] C. Mehring, U. Hehl, M. Kubo, M. Diesmann, and A. Aertsen, Biol Cybern **88**, 395 (2003), 6th Dynamic Brain Forum, Breisach, Germany, Sep 11-14, 2001.
[4] A. Kumar, S. Rotter, and A. Aertsen, J Neurosci **28**, 5268 (2008).
[5] M. van Rossum, G. Turrigiano, and S. Nelson, J Neurosci **22**, 1956 (2002).
[6] T. Vogels and L. Abbott, J Neurosci **25**, 10786 (2005).
[7] M. Diesmann, M. Gewaltig, and A. Aertsen, Nature **402**, 529 (1999).
[8] H. Cateau and T. Fukai, Neural Networks **14**, 675 (2001).
[9] W. Kistler and W. Gerstner, Neural Comput **14**, 987 (2002).
[10] V. Litvak, H. Sompolinsky, I. Segev, and M. Abeles, J Neurosci **23**, 3006 (2003).
[11] A. Aertsen, M. Diesmann, and M. Gewaltig, Journal of Physiology-Paris **90**, 243 (1996).
[12] M. Gewaltig, M. Diesmann, and A. Aertsen, Neural Networks **14**, 657 (2001).
[13] T. P. Vogels and L. F. Abbott, Nat Neurosci **12**, 483 (2009).
[14] A. Kepecs and G. Fishell, Nature **505**, 318 (2014).
[15] T. Klausberger and P. Somogyi, Science **321**, 53 (2008).
[16] J. S. Isaacson and M. Scanziani, Neuron **72**, 231 (2011).
[17] L. Roux and G. Buzsaki, Neuropharmacology **88**, 10 (2015).
[18] K. D. Harris and G. M. G. Shepherd, Nat Neurosci **18**, 170 (2015).
[19] Z. J. Huang, Neuron **83**, 1284 (2014).
[20] H. Taniguchi, Frontiers in Cellular Neuroscience **8** (2014), 10.3389/fncel.2014.00008.
[21] S. R. Olsen, D. S. Bortone, H. Adesnik, and M. Scanziani, Nature **483**, 47 (2012).
[22] D. Bortone, S. Olsen, and M. Scanziani, Neuron **82**, 474 (2014).
[23] S. X. Chen, A. N. Kim, A. J. Peters, and T. Komiyama, Nat Neurosci **18**, 1109 (2015).
[24] Y. Aviel, C. Mehring, M. Abeles, and D. Horn, Neural Comput **15**, 1321 (2003).
[25] T. Tetzlaff, M. Buschermohle, T. Geisel, and M. Diesmann, Neurocomputing **52-4**, 949 (2003).
[26] J. Kremkow, L. U. Perrinet, G. S. Masson, and A. Aertsen, J Comput Neurosci **28**, 579 (2010).
[27] A. Reyes, Nat Neurosci **6**, 593 (2003).
[28] K. Vincent, J. S. Tauskela, and J.-P. Thivierge, Frontiers in Computational Neuroscience **6** (2012), 10.3389/fncom.2012.00086.
[29] K. Nasrallah, R. A. Piskorowski, and V. Chevaleyre, eneuro **2(4)**, 1 (2015).
[30] K. Kohara, M. Pignatelli, A. J. Rivest, H.-Y. Jung, T. Kitamura, J. Suh, D. Frank, K. Kajikawa, N. Mise, Y. Obata, I. R. Wickersham, and S. Tonegawa, Nat Neurosci **17**, 269 (2014).
[31] M. Llorens-Martin, L. Blazquez-Llorca, R. Benavides-Piccione, A. Rabano, F. Hernandez, J. Avila, and J. DeFelipe, Frontiers in Neuroanatomy **8**, (2014).
[32] M. T. Schaub, Y. N. Billeh, C. A. Anastassiou, C. Koch, and M. Barahona, PLoS Computational Biology **11**, e1004196 (2015).
[33] S. Melzer, M. Michael, A. Caputi, M. Eliava, E. C. Fuchs, M. A. Whittington, and H. Monyer, Science **335**, 1506 (2012).
[34] A. Caputi, S. Melzer, M. Michael, and H. Monyer, *Macrocircuits*, Current Opinion in Neurobiology **23**, 179 (2013).
[35] G. Buzsaki, C. Geisler, D. Henze, and X. Wang, Trends in Neurosciences **27**, 186 (2004).
[36] A. Alonso and C. Kohler, Neuroscience Letters **31**, 209 (1982).
[37] A. T. Lee, D. Vogt, J. L. Rubenstein, and V. S. Sohal, The Journal of Neuroscience **34**, 11519 (2014).
[38] A. Alonso and C. Kohler, J. Comp. Neurol. **225**, 327 (1984).
[39] T. F. Freund and M. Antal, Nature **336**, 170 (1988).
[40] R. Tomioka, K. Okamoto, T. Furuta, F. Fujiyama, T. Iwasato, Y. Yanagawa, K. Obata, T. Kaneko, and N. Tamamaki, European Journal of Neuroscience **21**, 1587 (2005).
[41] S. Jinno, T. Klausberger, L. F. Marton, Y. Dalezios, J. D. B. Roberts, P. Fuentealba, E. A. Bushong, D. Henze, G. Buzsáki, and P. Somogyi, The Journal of Neuroscience **27**, 8790 (2007).
[42] A. I. Gulyas, N. Hajos, I. Katona, and T. F. Freund, European Journal of Neuroscience **17**, 1861 (2003).
[43] S. Pluta, A. Naka, J. Veit, G. Telian, L. Yao, R. Hakim, D. Taylor, and H. Adesnik, Nat Neurosci **18**, 1631

(2015).
[44] M. Karnani, J. Jackson, I. Ayzenshtat, J. Tucciarone, K. Manoocheri, W. Snider, and R. Yuste, Neuron **90**, 86 (2016).
[45] A. Litwin-Kumar and B. Doiron, Nat Neurosci **15**, 1498 (2012).